DOCUMENT RESUME

ED 171 536 SE 027 607

AUTHOR Schoenfeld, Alan H.
TITLE Presenting a Model of Mathematical Problem Solving.
PUB DATE Mar 79
NOTE 20p.; Paper presented at the annual meeting of the American Educational Research Association (San Francisco, California, April 8-12, 1979)

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *College Mathematics; *Course Descriptions; *Curriculum; Higher Education; *Instruction; *Problem Solving; *Research

ABSTRACT

Two questions are dealt with: (1) Can those strategies or behaviors which enable experts to solve problems well be characterized, and (2) Can students be trained to use such strategies? A problem-solving course for college students is described and the model on which the course is based is outlined in an attempt to answer these questions. The five-step model provides a guide to the problem-solving process. The stages are analysis, design, exploration, implementation, and verification. Pretest to posttest comparisons indicate that the students became more fluent at generating plausible approaches to problems. (MP)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
NATIONAL INSTITUTE OF EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
Alan H. Schoenfeld

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

PRESENTING A MODEL OF MATHEMATICAL PROBLEM SOLVING

by Alan H. Schoenfeld

ED171536

1. Background and Overview

In brief, my work in problem solving and this paper in particular deal with the following two questions:

1. Can we characterize those strategies or behaviors which enable experts to solve problems well, and
2. Having identified components of expert problem solving, can we train students to use some aspects of expert strategies, and thereby improve the students' problem solving?

My answer to both questions is a qualified and circumscribed "yes," but I should specify what I mean by "improve students' problem solving" before I elaborate. My standards (at least ideally) are rather strict. If a course is in general problem solving skills, then the students in the course should show a distinct improvement on a variety of problems not at all directly related to those in the course, when compared to a control group. In a general problem solving course, any problem at all, if the students have the appropriate background knowledge, is "fair game"; in a mathematics problem solving course for freshmen, any problem through freshman mathematics, or not requiring more background than freshman mathematics, is "fair game." Indeed, the testing should be done by someone who has no contact with the course save for a description of the backgrounds of the students and a ballpark idea of what "reasonable" problems are; the person(s) conducting the course should have no idea of the contents of the tests.

Under those conditions, one cannot be too sanguine about the success of a semester, or even a year-long course in "general problem solving". Little

SE 027 607

enough is known in detail about useful problem solving strategies, even in rather narrow problem solving domains, and less is known about how to teach them successfully. My decision, therefore, was to stay within one subject area -- mathematics. A number of "useful" strategies for mathematical problem solving have been described (mostly based on Polya's "heuristics"); yet, attempts to teach mathematical problem solving via heuristics have, generally, yielded rather unclear results. There are, in my opinion, a number of reasons that instruction in mathematical problem solving via general problem solving strategies has had only marginal success, among them:

1: the strategies have yet to be described in sufficient detail;

2: they are descriptive, rather than prescriptive; and

3: there are too many potentially useful strategies!

By "insufficient detail," I mean that the way we describe certain problem solving strategies are appropriate as a convenient label for a class of behaviors, but not detailed enough to specify how the labeled technique or strategy can be used. Consider the phrases "establish and exploit subgoals" and "consider special cases," two common heuristic strategies. How does one establish the subgoals in the first place? Having generated some plausible subgoals, how does one choose among them? Having reached a subgoal, does one exploit the method of solution, the result, or something else about it? And for problems like

a. Let $P(x)$ and $Q(x)$ be polynomials with "reversed" coefficients:

$$P(x) = \sum_{n=0}^{N} a_n x^n \qquad \text{and} \qquad Q(x) = \sum_{n=0}^{N} a_{N-n} x^n.$$

Determine and prove a relationship between the roots of $P(x)$ and $Q(x)$:

b. Let N be a positive integer. Determine the number of Divisors of N;

c. Show that in any circle, the central angle which cuts off a given arc is twice as large as the inscribed angle which cuts off the same arc;

the phrase "consider special cases" takes on rather different meanings. Frankly, I suspect that students will be unlikely to use the strategy on problems resembling these unless they have seen how to do so on these or similar problems.

Next we come to the distinction between descriptive and prescriptive strategies. Much of Polya's work, and that based on it, demonstrates the utility of heuristics by exemplifying their usage; that is, by showing how the heuristic approach results in rather elegant problem solutions. Yet this work rarely suggests why particular strategies were chosen, leading some to complain that in reading Polya's work one can only be a spectator to his tours de force without hope of being able to imitate them (Karplus, 1978). For a problem solving scheme to be useful, it must be prescriptive: that is, it must suggest how and when to use particular problem solving techniques.

Further, a long list of techniques, even if prescriptive, will be of little value unless it is embedded in some sort of manageable structure. The problem solver needs a means of narrowing down the collection of potentially useful strategies, and for effective budgeting of problem solving resources. As one such example, consider techniques of integration in first-year calculus. In indefinite integration, the individual techniques (integration by parts, by partial fractions, substitution, etc.) are nearly algorithmic and should present no major difficulties, but student performance is often worse than one would expect. This is because one must not only know how to use a particular technique; one must know that it is the "right" technique to use. In an experiment designed to test the utility of such global strategies (Schoenfeld, 1978), students taught a selection strategy significantly outperformed those who studied the "usual" way (working lots of problems) with less study time. In general problem solving, where there is a larger variety of potentially useful strategies and their application is far more subtle than in integration, the need for an overall organizational scheme is

all the much greater. Such a scheme is a model of "expert" problem solving in mathematics, based on observations of professional mathematicians solving a large number of problems. In my problem solving course, students are taught to follow the model of expert problem solving; in effect, I try to train them to think like experts. The course has been offered to upper division mathematics majors at U.C. Berkeley, and just this past term to lower division liberal arts students at Hamilton College. The model, the course, and the results will be described below.

2. The Model

The model outlined in figures 1 and 2 is meant to be both dynamic and prescriptive: all other factors being equal, the model provides a guide to the problem solving process. In the analysis stage of solving a problem (the first box of figure 1 and the first section of figure 2), the problem solver first faces the problem and grapples with it until (a) he or she has a "feel" for it, (b) the problem has been reformulated in a useful analytic way, and (c) the problem domain and the basic approach to the problem have been (temporarily) settled upon. In slightly more detail, the problem solver first reads (and rereads, if necessary) the problem and summarizes critical information contained in the problem statement including "given" and "goals." If appropriate, a diagram is drawn; special cases many be examined to exemplify the problem and see the range of possible or plausible answers; ballpark estimates of order of magnitude might be made. A mathematical context for the problem (say analytical geometry vs. Euclidean geometry) is established, and the problem may be re-cast in that context. (If one chooses to solve a problem using analytic geometry, the given terms such as "circle" may be translated

Schematic Outline of the Problem-Solving Strategy

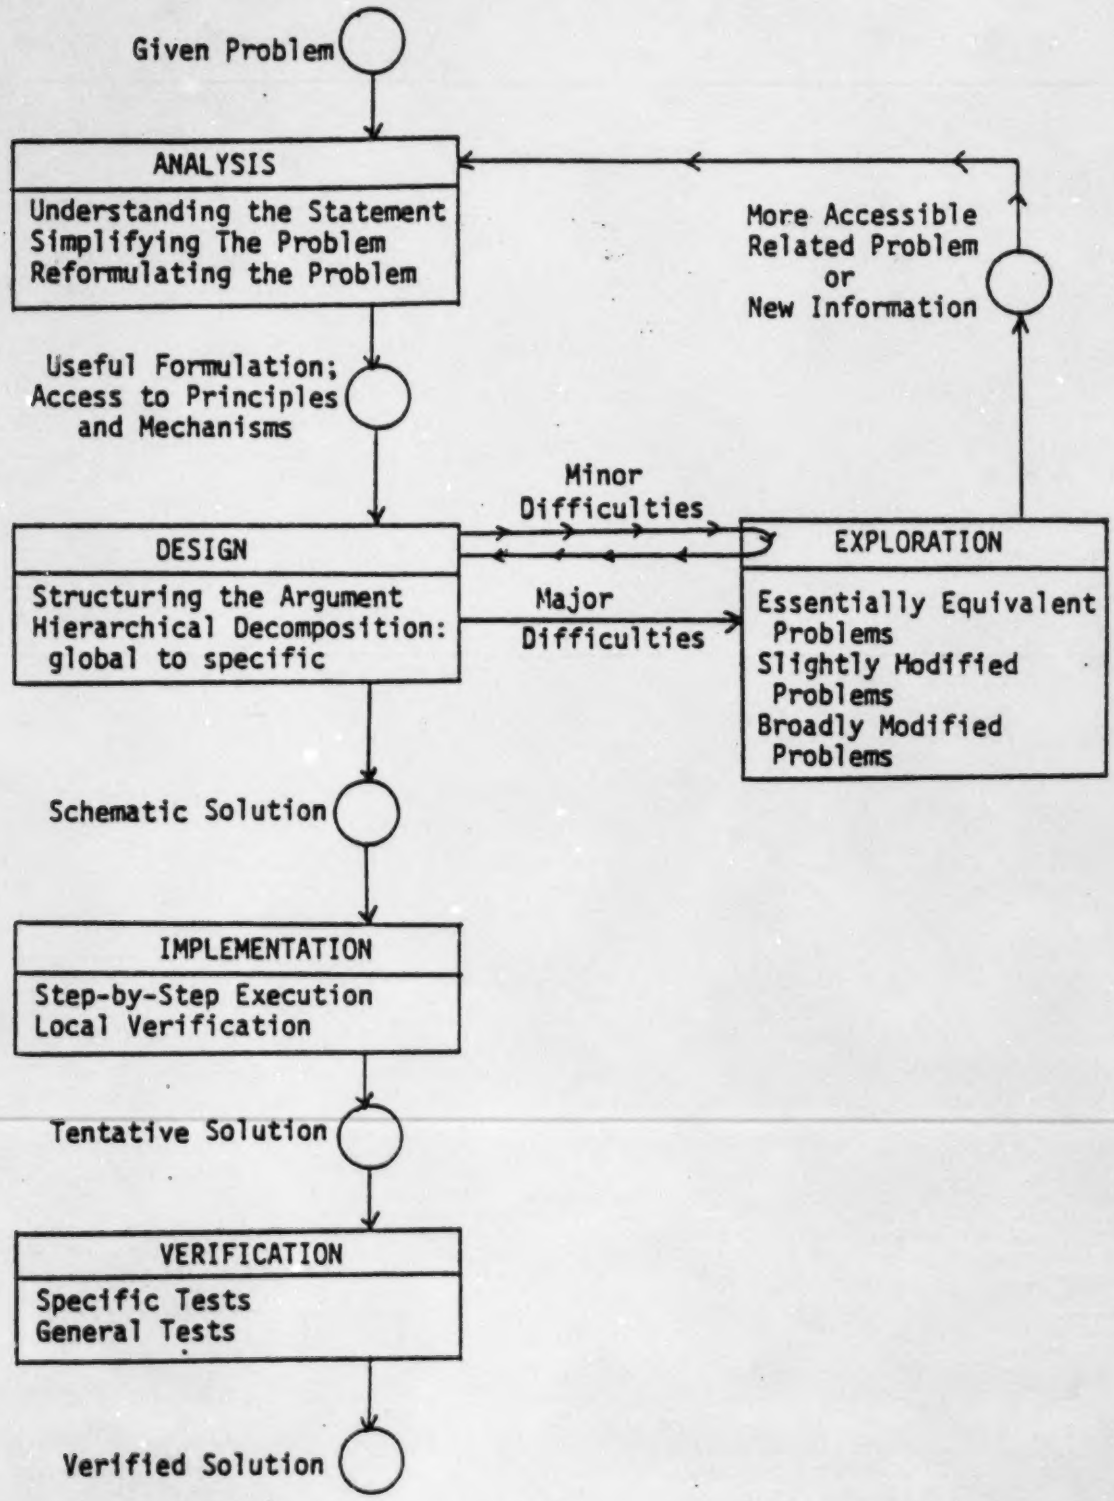

(figure 1)

FREQUENTLY USED HEURISTICS

ANALYSIS

1) DRAW A DIAGRAM if at all possible.

2) EXAMINE SPECIAL CASES:

   a) Choose special values to exemplify the problem and get a "feel" for it.

   b) Examine limiting cases to explore the range of possibilities.

   c) Set any integer parameters equal to 1, 2, 3,..., in sequence, and look for an inductive pattern.

3) TRY TO SIMPLIFY THE PROBLEM by

   a) exploiting symmetry, or

   b) "Without Loss of Generality" arguments (including scaling)

EXPLORATION

1) CONSIDER ESSENTIALLY EQUIVALENT PROBLEMS:

   a) Replacing conditions by equivalent ones.

   b) Re-combining the elements of the problem in different ways.

   c) Introduce auxiliary elements.

   d) Re-formulate the problem by

      i) change of perspective or notation

      ii) considering argument by contradiction or contrapositive

      iii) assuming you have a solution, and determining its properties

2) CONSIDER SLIGHTLY MODIFIED PROBLEMS:

   a) Choose subgoals (obtain partial fulfillment of the conditions)

   b) Relax a condition and then try to re-impose it.

   c) Decompose the domain of the problem and work on it case by case.

EXPLORATION (continued)

3) CONSIDER BROADLY MODIFIED PROBLEMS:

   a) Construct an analogous problem with fewer variables.

   b) Hold all but one variable fixed to determine that variable's impact.

   c) Try to exploit any related problem which have similar

      i) form

      ii) "givens"

      iii) conclusions.

Remember: when dealing with easier related problems, you should try to exploit both the RESULT and the METHOD OF SOLUTION on the given problem.

VERIFYING YOUR SOLUTION

1) DOES YOUR SOLUTION PASS THESE SPECIFIC TESTS:

   a) Does it use all the pertinent data?

   b) Does it conform to reasonable estimates or predictions?

   c) Does it withstand tests of symmetry, dimension analysis, or scaling?

2) DOES IT PASS THESE GENERAL TESTS?

   a) Can it be obtained differently?

   b) Can it be substantiated by special cases?

   c) Can it be reduced to known results?

   d) Can it be used to generate something you know?

(figure 2)

into equations, etc.). Finally attempts are made at preliminary simplifications, so that a "clean" and well-formulated version of the problem is ready for further study.

Properly construed, design is not localized into one box on the flow chart in figure 1, but is rather an "efficiency expert" whose role it is, at all points of the solution process, to ensure that the problem solver is using his or her time and energy wisely. In general, some principles to be adhered to in design are: (1) one should keep a global perspective -- at any point in a problem solution the problem solver should be able to say what is being pursued and why; what the options were; and what will be done with the results of the present operation, and (2) one should, unless there are strong reasons to do otherwise, proceed hierarchically. That is, solutions should be outlined first at a rough and qualitative level and elaborated in detail when warranted (that is, breadth before depth). As an example, the problem solver should not get involved in solving a messy set of equations until (i) alternatives which might make the solution unnecessary have been explored, (ii)it is clear that the solutions to the equations will help later on in the problem, and (iii) other parts of the solution have been elaborated to the point that it is clear that the energy spent in solving the equations will not be wasted.

As you can see in figure 2, exploration is divided into three stages. Generally, the suggestions in the first stage are either easier to employ than those in the second stage, or allow the problem solver to stay "closer" to the original problem; likewise for the relation between stages 2 and 3. Unless there are reasons to do otherwise, the problem solver in the exploration phase briefly considers those suggestions in stage 1 for plausibility, and then selects one or more and tries to exploit it. If the strategies in stage 1

prove insufficient, one proceeds to stage 2, and if need be, to stage 3. If substantial progress is made at any point in this process, the problem solver may either return to design to plan the balance of the solution or may re-enter analysis, with the expectation that the insights gained in exploration will allow for a "better" reformulation of the problem.

Implementation needs little comment. Verification, on the other hand, deserves more mention if only because it is so often slighted. At a local level, one can catch silly mistakes. At a global level, by reviewing the solution process one can often find alternative solutions, discover connections to other subject matter, and, on occasion , become consciously aware of useful aspects of the problem solution which can be incorporated into one's global strategy.

3. The Instruction

The Hamilton College academic calendar contains a three-and-a-half week long winter term, during which students enroll for one course only. The course enrollment this past year was 19; we expect similar enrollments (15-20) in coming years. Roughly half of the students are liberal arts majors, the rest potential science majors; most are freshmen or sophomores with one or three terms of college mathematics. Allowing time for testing and normal bureaucratic chores, this leaves fifteen or sixteen two-and-a-half hour classes devoted to problem solving; there is time for ample assignments between classes, with no other academic distractions. Thus, except for the brevity of the instructional period (three weeks offer little time for ideas to "sink in"), we have an unusually good environment for problem solving instruction.

In general, the course follows the model. At the beginning of the course the students are given the outlines in figures 1 and 2, along with a more

extended discussion of the model. On a "normal" day the class is given a handout with four or five problems; the students break into groups of three or four and work on the problems for a while. Then the class convenes as a whole, and solutions to some of the problems are discussed. As much as possible, the instructor tries to present solutions to problems within the context of the model. Where possible, "cues" are pointed out; the use of particular heuristic strategies is stressed; aspects of planning discussed, in detail (see below). In general, there is time for a discussion of perhaps half of the problems on the handout; the rest are made an assignment, to be discussed in a future class.

As a brief indication of the level of detail, let us consider the classroom discussion of the problem "construct, with straightedge and compass, the common external tangent to two circles." Some one suggested that the class consider an easier related problem, "construct the tangent to a circle from a point." The discussion of whether or not this was a good idea (if we found the construction, would it help solve the original problem?) led to a discussion in general of efficient ways to budget problem solving resources. To pose a related problem, or a subproblem, is to introduce an intermediate step between the givens and goals. This brings with it a series of questions and decisions: What is the likelihood of solving the intermediate problem, and how much work is involved? What is the likelihood of being able to use the solution towards the solution of the original problem, if you are able to solve it? How much information do you have about either? Which should you try first? For the above problem, we decided to see if the simpler problem could be exploited (it can), since someone remembered that it was a standard construction -- so that we could assumedly discover it, once we knew for certain that it would be useful.

4. The Results

The course in mathematical problem solving, based on the model outlined in section 2, has now been given twice: once at Berkeley in 1976, and this past winter at Hamilton College. More precisely, two very different courses based on the same model of expert performance have been given at Berkeley and at Hamilton. Given the complex web of skills required for "expert" performance in mathematics, I first offered the course at Berkeley for junior and senior undergraduate mathematics majors. The idea was that, while it might not be possible to condense a great deal of mathematical knowledge into a short course, it might well be possible to pull together for the students much that they had seen but had not yet recognized or codified for themselves. In a sense, then, I was saying "you have seen much of what I am about to show you, but not in coherent or organized fashion. In the normal course of events you would, over the next few years, discover that some of those strategies which you have used intermittently work rather consistently for you; you would, without being fully aware of it, come to use them more often. I shall single them out for you and make them explicit; being aware of them in this way, you can accelerate your development as a problem solver."

In large part, that is what happened (see "Can Heuristics be Taught?"). It was safe to assume that the students had a basic mastery of mathematical "tools of the trade," but, as early assignments and classroom sessions showed, little of the general strategic abilities of the expert. By the time of the final exam, the students were solving some problems they had been unable to approach coherently at the beginning of the term. Also, on part of the exam they were asked to indicate how they would approach a variety of problems if given ample time to work on them. Generally speaking,

they had learned to read some of the "cues" experts read in problems; these included deciding to approach certain problems by induction, by contradiction, by examining simpler analogous problems, etc.

Yet, despite these students' talents (respectable), predilections (mathematics majors, all) and backgrounds (upper division students), there were clear limits to what they could ingest. For example, problems which shared a similar deep structure but looked different on the surface were rarely recognized as being related (having used a sum of squares one way in one problem, the students failed to see that the next problem they were given, although much more complex in form, contained essentially a sum of squares and could be solved in the same way as the previous one). In fact, it might be most accurate to say that the students had learned a "first order" approximation to the model of expert performance. They knew, for example, that the presence of an integer parameter meant it might be appropriate to calculate a few cases and look for a pattern. Given the problem "How many subsets with an even number of elements does a set S with n elements have?" they would reliably calculate for n=1, 2, 3, 4 and draw the appropriate conclusion. But given the problem "how many divisors does the integer N have?" they would calculate a few cases, see no apparent pattern, and have no clear idea of how to proceed.

The situation at Hamilton was somewhat different. Due to a misunderstanding, "techniques of problem solving" was listed in the catalogue as a freshman level course. While this was frightening, it provided an opportunity to see if the more advanced backgrounds of the Berkeley students were truly necessary (as I had suspected), or whether it would be possible to teach aspects of the expert model to students whose mathematics backgrounds, for the most part, went no further than calculus. As an extra complicating variable, class size was 19, as opposed to the 8 I had taught at Berkeley.

In brief: fear was not justified, but concern was. The difference in level and background of the students had a tremendous impact on the running of the class. Topic coverage was changed, in that only problems solvable by the use of high school mathematics were used. But more than that, the depth of analysis and sophistication of the problems we discussed was lowered substantially. For example, consider the problem

d. Determine which numbers of the form aaa...a (the digit a is repeated n times) are perfect squares.

This was used as a problem on a take-home midterm in the Berkeley course (and the students did well on it); it was discussed in class, over the period of a week, in the more recent version of the course. Ostensibly, the reasoning needed to solve the problem is straightforward, and calls for nothing beyond high school mathematics. The observation that all perfect squares and in one of the digits 1, 4, 9, 6, 5, or 0 rules out strings of 2's, 3's, 7's, and 8's immediately (and a string of 0's is trivial, of course). A string of 5's is 5 times a string of ones; since a perfect square which has one factor of 5 must have at least two, and a string of ones does not have a factor of 5, no string of 5's is a perfect square. Likewise for factors of 6. Thus the only candidates are strings of 1's, 4's, and 9's. Factoring out 4's and 9's, the question becomes: when can a string of 1's be a perfect square? An investigation of all squares of the form $(50n + m)^2$, where m ranges from 1 to 50, shows that none of these end with even two 1's. Thus the only squares out of all the candidates are 0, 1, 4, and 9.

In my opinion, the difference between the two groups of students' performances on this problem was not a function either of general intelligence or of mathematical aptitude; as much as I could tell, there were no major differences between the two samples of students.

What was different, however, was the nature of the mathematics to which the students had been exposed. Through the calculus, the mathematics curriculum consists of what might be called didactic mathematics: students are shown well-defined techniques for solving certain classes of problems, and are expected to apply these techniques more or less directly. (This is much the same for using the quadratic formula in high school or solving max-min problems in calculus.) Thus the Hamilton students (for the most part) had not been exposed to typical patterns of abstract mathematical thought. It is these lines of thought, and the way they are used in solving problem (d), which make it hard. The argument is essentially inverted and negative; it proceeds by ruling out alternatives. Instead of asking "is this number a perfect square?" we ask "what properties do perfect squares have?" The realization that perfect squares end in certain digits mean that those that end in other digits are ruled out as candidates. Likewise, we eliminated strings of 5's and 6's through contradictions: because a perfect square must have an even number of factors of 5 and 6, and these did not, they were removed from consideration. Then, the realization that looking at the last digit did not provide enough information, causes a quick reversal: "what can I say about the last 2 digits of a perfect square?" The problem, then, is to find a representation which is convenient for doing so. Then the final argument is again negative: the list of perfect squares we obtain does not contain any which end in two 1's, so no string of 1's more than 1 digit long is a perfect square.

High school mathematics? Only ostensibly. For many of these students, the notion of mathematical proof is very hazy and the rationale for proving things even more so; they are not familiar with the paradigms of mathematical proof which more "advanced" students take for granted. This, as much as limitations on usable subject matter, determines what can and cannot

be done in such a course. "Problem solving" becomes more narrowly defined.

Students in the course were given matched pre- and post-tests. Since the course ended little over a month ago, much of the data has yet to be analyzed; the comments here are provisional and based on preliminary investigations. It is very clear that certain problem solving strategies can be learned reliably by the students, even over the short term. For example, the notion that an answer can be guessed empirically, and then verified was unfamiliar to the students. On the pre-test and post-test, respectively, were the problems

e. What is the sum of the first 89 odd numbers, and

f. What is the sum of the coefficients of $(x + 1)^{31}$?

Four students out of 19 saw that the sum of the odd numbers gives 1, 4, 9, 16, ..., n^2 respectively, so that the answer to (e) is 89^2; another six paired terms at the beginning and end of the sequence (Gauss' argument) to get the answer. Four of those explicitly remembered having seen the Gauss-type argument before. On the post-test, 16 of the 19 students solved the problem completely. In the course, no problem even vaguely resembling (f) in surface structure was discussed.

There were comparab'e results for two problems which could be argued by contradiction. Seven of the students thought to argue by contradiction on the pre-test; on the post-test, 14 did -- although the post-test problem was substantially harder than the pre-test problem, and only eight solved it.

In general, where specific "clues" were present in problem structure (such as in problems which can be solved by induction, or "fewer variables") the students recognized the clues on the final exam and did significantly better on it than they had on the pre-test. For other problems, where it is hard <u>a priori</u> to single out a "proper" approach, the results are less clear.

Pre-test to post-test comparisons indicate that the students became more fluent at generating plausible approaches to problems; we have the students' subjective measures of their performance on the matched problems, and the objective scoring to back this up. Statistical tests have not yet been run to see if their performance is significantly improved on the more difficult problems; I suspect that the results on those problems alone will not be significant. In fact, it is in dealing with general strategies -- to generate and evaluate plausible approaches to a problem, and to budget one's resources efficiently -- that the course needs most to be revised. Much of the three weeks of intensive problem solving was spent providing the students with the tools they needed; the cost of this was that I could not spend enough time on broader issues. To some degree, this was a pilot study of both the instruction and the test instruments. To pass on a quote passed on to me by Jill Larkin; "First courses, like first pancakes, rarely turn out right." The proportions in the recipe for the course will be changed somewhat, but the ingredients will remain pretty much the same. The course clearly had an impact, reflected in the test data analyzed so far and in course evaluations; the next time around I hope for more.

References

Brown, J.S. Diagnostic Models for Procedural Bugs in Basic Mathematical Skills, Report 3669, Bolt, Beranek & Newman, Cambridge, 1977.

Brown, J.S. Artificial Intelligence and Learning Strategies, Report 3634, Bolt, Beranek & Newman, Cambridge, 1977.

Campbell, D., & Stanley, J. Experimental and Quasi-Experimental Designs for Research. Chicago: Rand-McNally, 1963.

Goldberg, D.J. The effects of training in heuristic methods on the ability to write proofs in number theory (Doctoral Dissertation, Columbia University, 1974). Dissertation Abstracts International, 1974, 4989-B. (University Microfilms N. 75-7836)

Hayes-Roth, B, & Hayes-Roth, F. Cognitive Processes in Planning, Working Note for the O.N.R.; Santa Monica: Rand Corporation, 1978.

Kantowski, M.G. Processes involved in mathematical problem solving. Journal For Research in Mathematics Education, May, 1977, 163-180.

Karplus, R.A. A review of Polya's Mathematical Methods in Science, in the American Mathematical Monthly, April, 1978.

Kilpatrick, J. Variables and Methodologies in Research in Problem Solving, Paper for research workshop on Problem Solving in Mathematics education, University of Georgia, 1975.

Klahr, D. Cognition and Instruction. New Jersey: Lawrence Earlbaum Associates.

Kleinmuntz, B. Problem Solving: Research, Method, and Theory. New York: Wiley, 1966.

Landa, L.M. The ability to think: can it be taught? Soviet Education, Vol. XVIII, 5, March 1976.

Larkin, J. Understanding Relations in Physics Unpublished Doctoral Dissertation, University of California, Berkeley, 1974.

Larkin, J. Skilled Problem Solving in Physics: A Hierarchical Planning Model Berkeley: University of California, Berkeley, 198.

Lucas, J.F. An exploratory study on the diagnostic teaching of heuristic problem solving strategies in calculus (Doctoral dissertation, University of Wisconsin, 1972). Dissertation Abstracts International, 1972, 6825-A. (University Microfilms 72-15, 368)

Miller, G., Galanter, E. and Pribram, D. Plans and the Structure of Behavior, New York: Holt, Reinhard & Winston, 1960.

Polya, G. How to Solve It. New Jersey: Princeton University Press, 1945.

Polya, G. Mathematical Discovery, (Vols. 1 & 2). New York: Wiley, 1962 & 1965.

Reif, F., Editor. Proceedings of the Carnegie-Mellon Conference on Problem Solving and Education. New Jersey: L.E.A., 1979.

Rubenstein, M. Patterns of Problem Solving. New Jersey: Prentice-Hall, 1975.

Schank, R. & Abelson, R. Scripts, Plans, Goals, and Understanding. New Jersey: Lawrence Earlbaum Associates, 1977.

Schoenfeld, A. Presenting a Strategy for Indefinite Integration. American Mathematical Monthly, October 1978.

Schoenfeld, A. Can Heuristics be Taught? In Cognitive Process Instruction, J. Lochhead, Editor. Philadelphia: Franklin Institute Press, 1979.

Schoenfeld, A. Heuristic Variables in Instruction, Chapter 10 of Task Variables in Mathematical Problem Solving, Edited by G. Goldin and E. McClintock, ERIC, 1979.

Schoenfeld, A. About Heuristics: What they Are and How we Can Teach Them. In the 1980 Yearbook of the N.C.T.M.

Simon, D. and Simon, H. Individual Differences in Solving Physics Problems C.I.P. Working Paper 342, Pittsburgh, Carnegie-Mellon University, 1977.

Smith, J.R. The effects of general versus specific heuristics in mathematical problem solving tasks (Doctoral dissertation, Columbia University, 1973). Dissertation Abstracts International, 1973, 2400-A. (University Microfilm No. 73-26, 637)

Webb, N.L. An Exploration of Mathematical Problem Solving Processes Unpublished doctoral dissertation (Stanford University, 1975)

Wickelgren, W. How to Solve Problems. San Francisco: W. H. Freeman, 1975.

Wilson, J.W. Generality of heuristics as an instructional variable (Doctoral Dissertation, Stanford University, 1967). Dissertation Abstracts, 1967, 2575-A. (University Microfilms 67-17, 526)